



(11)

EP 4 080 501 A1

(12)

EUROPEAN PATENT APPLICATION
published in accordance with Art. 153(4) EPC

(43) Date of publication:
26.10.2022 Bulletin 2022/43

(51) International Patent Classification (IPC):
G10L 17/26^(2013.01) G10L 25/30^(2013.01)

(21) Application number: **19870072.6**

(52) Cooperative Patent Classification (CPC):
G10L 17/26; G10L 25/30

(22) Date of filing: **16.12.2019**

(86) International application number:
PCT/ES2019/070849

(87) International publication number:
WO 2021/123462 (24.06.2021 Gazette 2021/25)

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR
Designated Extension States:
BA ME
Designated Validation States:
KH MA MD TN

(72) Inventors:
• **PERERO CODOSERO, Juan Manuel**
28100 Alcobendas (ES)
• **ESPINOZA CUADROS, Fernando**
28100 Alcobendas (ES)

(71) Applicant: **SIGMA.AI SL**
28050 Madrid (ES)

(74) Representative: **Elzaburu S.L.P.**
Edificio Torre de Cristal
Paseo de la Castellana 259 C, planta 28
28046 Madrid (ES)

(54) **METHOD AND SYSTEM TO ESTIMATE SPEAKER CHARACTERISTICS ON-THE-FLY FOR UNKNOWN SPEAKER WITH HIGH ACCURACY AND LOW LATENCY**

(57) A computer-implemented technique is presented for profiling an unknown speaker. A DNN-based frame selection allows the system to select the relevant frames necessary to provide a reliable speaker characteristic estimation. A frame selection module selects those frames that contain relevant information for estimating a given speaker characteristic and thereby contributes to the accuracy and the low latency of the system. Real-time speaker characteristics estimation allows the system to estimate the speaker characteristics from a speech segment of accumulated selected frames at any given time. The frame level processing contributes to the low latency as it is not necessary to wait for the whole speech utterance to predict a speaker characteristic but rather a speaker characteristic is estimated from only a few reliable frames. Different stopping criteria also contribute to the accuracy and the low latency of the system.




Figure 2

EP 4 080 501 A1

Description

FIELD

[0001] The present disclosure relates to speaker profiling by means of speaker characteristic estimation from the speech signal. The speech signal contains specific characteristic information related to physical (e.g., gender, age, height, weight, BMI), geographical (e.g., language, accent), sociolect (e.g., dialect) and/or ethnic background (e.g., race) of the speaker.

BACKGROUND

[0002] Speech analytics technologies have been applied to several areas like communication, marketing, and security. Speech analytics adds value to these areas in different ways by providing specific information about speakers. This information can be obtained from the speech signal. The speech signal contains speaker characteristics which can be extracted automatically to know specific aspects of a person.

[0003] There is a strong benefit to perform the extraction of speaker characteristics in real-time, e.g., for security alerts it is necessary to have a minimum response time, or for sales agents working in call centers it is very helpful to make recommendations personalized to a particular speaker in "conversation time". To get quick responses, a system low latency is required.

[0004] System accuracy is also critical for these applications because it is important to make the right decision, as mistakes can have negative consequences for the security of the people involved, or provide a bad experience to customers if the response does not fulfill their expectations.

[0005] A few prior art systems are noted below. One system focuses on identifying speakers for call centers. Call center applications need to work in real-time to provide an output in a reasonable time interval. However, this system does not provide information about speaker characteristics. It only provides the identity of known speakers, i.e. speakers previously enrolled in the system. Moreover, the implementation is not based on DNN architectures.

[0006] Other speaker identification systems operating in real-time for applications such as web-based conferencing have emerged as well. Most of these speaker identification/recognition applications have been used to detect who is speaking along the conversation, also called speaker diarization, or generating a profile (e.g., age, gender, job occupation) about who is speaking with the goal reinforcing the speaker identification task with additional speaker characteristics to overcome the lack of context information, or associate additional speaker data to a speaker that was previously enrolled. Unlike these approaches, the system proposed in this disclosure performs speech analytics for unknown speakers (i.e. not enrolled speakers in the system). The proposed

system provides characteristics 'on-the-fly' for an unknown speaker like the gender, age, height, body mass index (BMI), or language from the speech signal.

[0007] Speaker recognition systems have also been used to train hierarchical speaker clusters. Voice data from known speakers and associated label are employed in the cluster training process. Labels contain information about gender, age and native language. Clusters are based on Gaussian Mixture Models (GMMs) and Mel-Frequency Cepstral Coefficients (MFCCs) were used in the training process. During the recognition phase a test speaker is associated to a speaker identity and at the same time is assigned to a speaker cluster. In contrast, the proposed system is based on Deep Neural Networks (DNNs) models, one model for each characteristic. These models are trained using speech from speakers who are different from the test ones (unknown speakers).

[0008] The technique presented in this disclosure makes use of frame selection to extract the relevant frames of interest. While this disclosure proposes the use of DNNs to extract the best frames of interest for each speaker characteristic extracted, other techniques such as Voice Activity Detection (VAD) could potentially be used to separate speech for silence, so that only speech segments would be considered in the estimation of speaker characteristics. However, while VAD makes use of a short-term energy analysis, Deep Learning-based analysis was shown, in the context of a speaker verification task, to improve performance to carry out frame selection.

[0009] Other known techniques use physical characteristics to improve the accuracy of the speaker identification. It consists in the extraction and analysis of vocal characteristics and frequently used words in the voice segment. The extracted information helps to fill a database linking the age, gender even ethnic group with the speaker identity. Similarly, techniques have been developed to provide a likely age range associated with the speaker based on the evaluation of a spoken phrase in terms of prosody such as speaker rate, number and length of pauses in the spoken phrase, misarticulation of a portion of the spoken phrase, and vocal intensity of the caller's voice.

[0010] Regarding the topic of speaker characteristics prediction, there is work in the art related to age prediction, height estimation or BMI classification from speech. None of these methods need to obtain live results and they have not been developed using DNN-based algorithms. In contrast, the proposed system achieves speaker characteristics 'on-the-fly' and is based on DNN architectures.

[0011] Breakthroughs in Deep Learning has led to development of accurate methods estimating speaker characteristics offline with DNN-based models (e.g., age estimation) and by mixing DNN-models with other technologies such as facial recognition. In contrast, the proposed system, which is also based on DNNs architectures, performs live speaker characteristics estimation providing

results with a very low latency.

[0012] Low latency is one of the main points of this disclosure. Achieving the extraction of speaker characteristics in the shortest amount of time allows to provide better services by enhancing the user experience. Such systems are able to help sales agents working in call centers or intelligent agents (or bots) making recommendations to offer a customized product or specific services according to the customer characteristics. Using "on-the-fly" specific characteristics extraction enables systems to offer new products, services or specific content to customers nearly real-time. Just a few seconds can make a difference between a high-quality service and a poor service. Another important point to keep in mind is the accuracy. The higher the accuracy, the better the services will be.

[0013] One of the main challenges to develop the extraction of speaker characteristics "on-the-fly" with high accuracy is related to the selection of the speech features to use and to know when a specific speaker characteristic has been extracted so that a real-time action based on the extraction of the speaker characteristic can take place. These challenges have been solved in this disclosure.

[0014] This section provides background information related to the present disclosure which is not necessarily prior art.

SUMMARY

[0015] This section provides a general summary of the disclosure, and is not a comprehensive disclosure of its full scope or all of its features.

[0016] In one aspect, a computer-implemented method is provided for profiling an unknown speaker, such that a profile for a speaker includes a set of speaker characteristics. The method includes: receiving a speech signal from an unknown speaker, where the speech signal is comprised of a plurality of data frames; for each speaker characteristic in the set of speaker characteristics, selecting a subset of data frames from a corresponding speech signal in the plurality of speech signals and thereby exclude other data frames in the corresponding speech signal from the subset of data frames, where the data frames in the subset of data frames contain information about the speaker characteristic of interest; and for a given speaker characteristic, generating a feature vector from the subset of selected data frames corresponding to the given speaker characteristic and predicting a value for the given speaker characteristic from the feature vector using machine learning. The speaker characteristics may include but are not limited to age, height, body mass index, gender, accent, dialect and race.

[0017] The speech signal may be captured from the unknown speaker using a microphone.

[0018] In one embodiment, selecting a subset of data frames from a speech signal is implemented using at least one classifier (e.g., a neural network) for each

speaker characteristic.

[0019] Similarly, predicting a value for the given speaker characteristic may be implemented using a neural network (e.g., a feed-forward artificial neural network) for each speaker characteristic.

[0020] The method may further include assessing reliability of the predicted value for the given speaker characteristic; and adding the predicted value for the given speaker characteristic to a profile for the unknown speaker when the predicted value for the given speaker characteristic is deemed reliable. In one embodiment, reliability of the predicted value for the given speaker characteristic is assessed by comparing the predicted value to a prediction interval, and deeming the predicted value reliable when the predicted value falls within the prediction interval, where the prediction interval is based on a variance of a difference between predicted values for the given speaker characteristic and true values for the given speaker characteristic.

[0021] In another aspect, a computer-implemented method is provided for profiling an unknown speaker, such that a profile for a speaker includes a set of speaker characteristics. The method includes: a) receiving a speech signal from an unknown speaker, where the speech signal is comprised of a plurality of data frames; b) accumulating data frames from the plurality of data frames in a buffer; c) generating, at a given time, a feature vector from the data frames accumulated in the data buffer; d) predicting a value for a given speaker characteristic from the feature vector using machine learning, and assessing reliability of the predicted value for the given speaker characteristic; e) adding the predicted value for the given speaker characteristic to a profile for the unknown speaker when the predicted value for the given speaker characteristic is deemed reliable; and repeating steps b) - d) when the predicted value for the given speaker characteristic is deemed unreliable.

[0022] The speech signal may be captured from the unknown speaker using a microphone.

[0023] In one embodiment, the reliability of the predicted value for the given speaker characteristic is assessed by comparing the predicted value to a prediction interval, and deeming the predicted value reliable when the predicted value falls within the prediction interval, where the prediction interval quantifies uncertainty related to the predicted value, for example based on a variance of a difference between predicted values for the given speaker characteristic and true values for the given speaker characteristic.

[0024] In another embodiment, the reliability of the predicted value for the given speaker characteristic is assessed by comparing the predicted value to a prediction interval, where the prediction interval is based on a variance of a difference between predicted values for the given speaker characteristic and true values for the given speaker characteristic; computing a weighted average using the predicted value and a previous weighted average when the predicted value falls within the prediction

interval, where the previous weighted average is computed using at least one predicted value for the given speaker characteristic predicted from a feature vector generated at a time earlier than the given time; comparing the weighted average to another prediction interval; and deeming the predicted value reliable when the weighted average falls within the prediction interval.

[0025] For each speaker characteristic in the set of speaker characteristics, a value for a particular speaker characteristic is predicted from the feature vector using a corresponding neural network, where a different neural network is implemented for each speaker characteristic in the set of speaker characteristics.

[0026] The steps of accumulating data frames from the plurality of data frames in a buffer, generating a feature vector from the data frames accumulated in the data buffer, and predicting a value for each speaker characteristic in the set of speaker characteristics that is deemed unreliable are repeated until each speaker characteristic in the set of speaker characteristics is deemed reliable.

[0027] In some embodiments, the feature vector is generated from a number of data frames which varies over time.

[0028] In yet another aspect, a system is presented for profiling an unknown speaker. The system includes: a frame extractor, a feature extractor, a speaker estimator and a speaker profiler. The frame extractor is configured to receive a speech signal from an unknown speaker and accumulate data frames from the speech signal in a buffer. The feature extractor is configured to generate a feature vector from the data frames in the buffer. The speaker estimator is configured to receive the feature vector from the feature extractor and, for each speaker characteristic in the set of speaker characteristics, predict a value for a speaker characteristic from the feature vector using a first set of neural networks, where the speaker estimator implements a different neural network for each speaker characteristic in the set of speaker characteristics. The speaker profiler is configured to receive the predicted values for each speaker characteristic from the speaker estimator, where the speaker profiler assesses the reliability of the predicted value for each speaker characteristic and adds the predicted value for a given speaker characteristic to a speaker profile when the predicted value for the given speaker characteristic is deemed reliable. A microphone is configured to capture the speech signal from the unknown speaker.

[0029] The system may further include a frame selector configured to receive the speech signal from the frame extractor. The frame selector selects a subset of data frames from a corresponding speech signal in the plurality of speech signals and accumulates the subset of data frames in a buffer corresponding to the speaker characteristic, thereby excluding other data frames in the corresponding speech signal from the subset of data frames, where the data frames in the subset of data frames contain information about the speaker characteristic of interest. The frame selector may select a subset of data

frames from each of the plurality of speech signals using a second set of neural networks, where the frame selector implements a different neural network for each speaker characteristic in the set of speaker characteristics.

[0030] In one embodiment, the speaker profiler assesses reliability of the predicted value for a given speaker characteristic by comparing the predicted value to a prediction interval and deeming the predicted value reliable when the predicted value falls within the prediction interval, where the prediction interval quantifies uncertainty related to the predicted value.

[0031] In other embodiments, the speaker profiler assesses reliability of the predicted value by comparing the predicted value to a prediction interval, where the prediction interval is based on a variance of a difference between predicted values for the given speaker characteristic and true values for the given speaker characteristic; computing a weighted average using the predicted value and a previous weighted average when the predicted value falls within the prediction interval, where the previous weighted average is computed using at least one predicted value for the given speaker characteristic predicted from a feature vector generated at a time earlier than the given time; comparing the weighted average to another prediction interval; and deeming the predicted value reliable when the weighted average falls within the prediction interval.

[0032] Further areas of applicability will become apparent from the description provided herein. The description and specific examples in this summary are intended for purposes of illustration only and are not intended to limit the scope of the present disclosure.

DRAWINGS

[0033] The drawings described herein are for illustrative purposes only of selected embodiments and not all possible implementations, and are not intended to limit the scope of the present disclosure.

Figure 1 is a diagram showing a proposed system for "on-the-fly" speaker characteristics estimation and user profiling.

Figure 2 is a diagram explaining the frame selection; the diagram covers the steps from when a live speech signal is received to the one when a frame has been selected or not for each characteristic.

Figure 3A is a diagram showing the training data used for the frame selection process and a module to select a threshold for reliable-frame criterion, where the training data is used as input for each DNN-based classifier.

Figure 3B is a graph depicting an example of frame confidence score distribution for the age characteristic to calculate the threshold for the reliable-frame

selection criterion.

Figure 4 is a diagram showing an example of the test process at the end of the frame selection module.

Figure 5 is a diagram explaining the live feature extraction workflow from the first step when a vector of coefficients has been extracted to the last one when the vector is ready to be used by the speaker characteristics estimation module.

Figure 6 is a diagram showing the training data as input for each DNN-based classifier, along with different values or classes used for each characteristic.

Figure 7 is a diagram showing an example of the test process for the gender and age characteristics.

Figure 8 is a diagram showing the workflow for the stopping criteria from the first step, when a minimum number of frames criterion has to be satisfied to the last step when a reliable prediction is provided for each characteristic.

Figure 9 is a diagram illustrating how the prediction interval (PI) estimation used for determining the minimum number of accumulated frames is estimated at training time and how the test phase is conducted, including how the prediction interval (PI) estimation process is based on the variance of the difference between the true value and the predicted values (top) and how the evaluation of a speech segment of accumulated selected frames at time t by the minimum accumulated frame number criterion (bottom).

Figure 10 shows the standard Normal z-table for two-tailed probabilities, where the table represents the area or percentage under the standard normal curve corresponding to any z-score.

Figure 11 is a graph showing the age prediction system performance as a function of the number of available frames.

Figure 12 is a diagram illustrating the prediction interval (PI) estimation for the reliable prediction criterion and the evaluation of a speech segment by the aforementioned criterion, where the prediction interval (PI) process is based on the variance of the difference between the true value and the predicted values (top) and the evaluation of the reliable prediction criterion (bottom).

Figure 13 is a diagram illustrating the threshold's estimation for each class from Log-Likelihood Ratios (LLRs) for speaker characteristics with categorical values and the minimum number of frames criterion, where the thresholds' estimation is based on the

LLRs (top) and the evaluation of the minimum number of frames criterion during the test phase (bottom).

Figure 14 is a diagram illustrating the threshold's estimation for each class from LLRs and the reliability prediction criterion, where the thresholds' estimation is based on the weighted average LLRs using a training dataset (top) and the evaluation of the reliable prediction criterion during the test phase (bottom).

[0034] Corresponding reference numerals indicate corresponding parts throughout the several views of the drawings.

DETAILED DESCRIPTION

[0035] Example embodiments will now be described more fully with reference to the accompanying drawings.

[0036] Figure 1 shows the general system to estimate 'on-the-fly' speaker characteristics and user profiling. In a nutshell, the speaker characteristics prediction is performed at each time t . The system will estimate those characteristics using a segment of N selected frames accumulated at a time t . As a result, the speaker characteristics estimation is performed in real time (or nearly real-time).

[0037] The system follows a novel approach for speaker characteristics prediction. It performs the predictions using short-term cepstral features modeling, that is cepstral vectors representing the spectral content of speech window modeled through a generative model such as Gaussian Mixture Model (GMM), and a DNN-based prediction model. As mentioned, the unique capabilities of the disclosure as compared to prior art systems lie in the low latency and high accuracy of the speaker characteristics extraction process.

[0038] The disclosure includes a frame selection module 010, which is one novel component of the disclosure, contributing to the high accuracy and low latency, a feature extraction module 020, a speaker characteristics estimation module 030 based on DNNs, and novel stopping decision criteria 040 to provide accurate predictions of speaker characteristics. As used herein, the term module may refer to, be part of, or include an Application Specific Integrated Circuit (ASIC), an electronic circuit, a computer processor that execute one or more software or firmware programs, a combinational logic circuit, and/or other suitable components that provide the described functionality.

[0039] Figure 2 shows the frame selection 010 workflow. For each specific characteristic, the frame selection module 010 aims at filtering frames which do not contain relevant information to extract the specific speaker characteristic. As a result, the frames containing very little information to identify a specific user characteristic will be eliminated and not be processed. This frame selection process contributes to a faster convergence towards

identifying a specific speaker characteristic. Another consequence of processing only relevant frames is an improvement in the accuracy of the speaker characteristics prediction.

[0040] The frame selection module 010 consists of three sub-modules: frame extraction 011, frame confidence score estimation 012, and reliable-frame criterion 013. In the example embodiment, the voice is collected at frame level by means of an independent microphone or one embedded in an electronic device.

[0041] For frame confidence score estimation 012, the confidence scores are calculated for each frame. This confidence score will be used to select the most suitable frames to predict speaker characteristics. In one embodiment, this block is designed with several DNN-based classifiers in parallel. For example, one classifier is implemented per speaker characteristic. Every classifier may be a Convolutional Neural Networks (CNNs) implementation.

[0042] Figure 3 explains how the classifiers are trained and the probabilities of each class, i.e. confidence scores, are obtained.

[0043] Figure 3 on the left shows the training data provided as input to each DNN-based classifier. Each classifier is trained with the speech frames labeled for a specific characteristic. In this example, gender and age are the speaker characteristics used.

[0044] The training dataset consists of voice samples and its associated labels per characteristic. Voice samples, which have a duration of 25 ms, are raw speech data to which a Hamming window is applied with a shift of 10 ms to obtain the next frame at the subsequent iteration.

[0045] The labels corresponding to the characteristics expressed as numerical values, e.g., age, height and BMI, are mapped to a number of classes per characteristic, $K_{characteristic}$. Each $K_{characteristic}$ value will be chosen according to the required levels of granularity.

[0046] For the labels corresponding to categorical values (e.g., gender and language), the number of classes, $K_{characteristic}$, will depend on the number of classes which are in the database. For example, $K_{language}$ value will depend on the number of languages represented in the database.

[0047] Figure 3 includes example K values for gender and age. In case of gender, $K_{gender} = 2$ is used to classify male speech and female speech. In case of age, $K_{age} = 3$ is defined after deciding to classify the age in the classes: youngs, middle-age and elders. The same process would apply to the remaining characteristics of interest.

[0048] Throughout the training process a class balanced database is used to train appropriately every class for each characteristic.

[0049] In the example embodiment, the architecture of the classifiers is the same for all the characteristics. It is composed of several filter stages followed by a classification stage. The initial filter stages are also known as feature extraction phases. The architecture consists of

convolutional layers, max-pooling layers, and non-linear activation functions. The classification stage consists of a fully-connected layer and a softmax layer. The softmax layer is used to represent the probability distribution of each class i for each frame x , $p(i|x)$.

[0050] More specifically, the softmax function is as follows:

$$\sigma(z)_i = \frac{e^{z_i}}{\sum_{j=1}^K e^{z_j}}$$

$i = 1, \dots, K$; $z = (z_1, \dots, z_K) \in \mathbb{R}^K$ where the standard exponential function is applied to each element z_i of the vector z . The values $\sigma(z)_i$, which represent the probabilities for each class, are normalized to ensure that the classifier outputs are between 0 and 1. The sum of the components $\sigma(z)$ is 1. The probabilities of the classes represent the confidence scores of those classes.

[0051] Once the classifiers are trained, the confidence scores are provided as input to the reliable frame selection module 013. In one example embodiment, each characteristic classifier has a specific threshold associated ($\gamma_{characteristic}$). It is used to determine when a frame is reliable.

[0052] To determine the threshold for a specific characteristic, first run the previous classification step for a specific dataset using the frame confidence score estimation module. Then, use the classification outputs, i.e. confidence scores, to calculate the confidence score distribution.

[0053] Figure 3 on the right depicts an example of a frame confidence score distribution for the age characteristic. Only the correctly classified frames, i.e. when the most probable class is the right class, are used according to the label in the dataset. The distribution is used to determine the threshold as the reliable-frame criterion 013. In one example, a third quartile Q_3 criterion is used to define the threshold.

[0054] Third quartile is also known as the 75th percentile. The result is a value under which we found the 75% of the correctly classified samples of the distribution. This value will be considered the threshold. In this example, the threshold γ_{age} , is 0.8. For future input frames when the confidence score will be greater than γ_{age} , the classification is considered to be correct. This process will be carried out for all the characteristics of interest (e.g., gender, height, BMI, language).

[0055] During the test process for each characteristic, the system will select the frames that are relevant for the characteristic of interest based on the reliable-frame selection criterion 013. Each input frame will be evaluated for each characteristic as a reliable frame or not.

[0056] Figure 4 shows an example of the test process for gender and age. One frame will get different confidence scores depending on the characteristic. In the ex-

ample shown, for the gender characteristic the most probable class gets a confidence score of 0.91 (corresponding to class $i=1$ or "Female"). For the age characteristic, the most probable class gets a confidence score of 0.75 for the class $i=2$ ("Middle Age").

[0057] If one assumes that the threshold for gender $\gamma_{\text{gender}} = 0.85$ and the threshold for age $\gamma_{\text{age}} = 0.80$, the input frame is considered to be reliable for gender, but not for age since its associated confidence is lower than the age threshold.

[0058] The above explanation illustrates that one frame can be more representative of a specific characteristic as compared to others. This mechanism contributes to reducing the latency as some characteristics will be predicted before others because a sufficient number of frames have been selected allowing a reliable characteristic to be extracted. In other words, all the speaker characteristics are not extracted at the same time.

[0059] After the frame selection module 010, selected frames are accumulated in a buffer into a speech segment at time t , and this segment feeds the live feature extraction module 020.

[0060] Figure 5 shows the live feature extraction module 020 workflow. The live feature extraction module 020 consists of three sub-modules: feature vector extraction 021, feature vector modeling 022, and class separation maximization 023.

[0061] For feature vector extraction 021, 40-dimensional Mel Cepstrum coefficients (MFCCs) vectors are used in an example embodiment. Hamming windows (e.g., 25 ms) are slid by 10 ms at each iteration. Then a Cepstral Mean Variance Normalization (CMVN) is applied. It should be noted that other types of feature vector (e.g., Linear Predictive Coding, Cepstrum coefficients or Fast Fourier Transform coefficients) can be used.

[0062] For features vector modeling 022, an i -vector extractor module based on a gender-independent UBM-GMM of 512 components is used with a total variability subspace of dimensionality 400. It should be noted that other feature vector modeling algorithms can be used, such as DNN-based embeddings.

[0063] For class separation maximization 023, a Linear Discriminant Analysis (LDA) is used. LDA applies a dimensionality reduction to 200 components, and projects the orthogonal axes to a new space to better discriminate between different classes.

[0064] The output of the live feature extraction module 020, which are i -vectors projected in a smaller dimensional space by the LDA, is used as input to the speaker characteristics estimation module 030.

[0065] The speaker characteristics estimation module 030 is designed with several DNN-based models in parallel. In an example embodiment, one model per speaker characteristic, e.g., gender, age, height, BMI, language, is implemented. Every model consists of a four layer, fully-connected DNN.

[0066] Figure 6 shows the training data used in the training process. In this example, gender and age are

the speaker characteristics used.

[0067] The training dataset consists of voice samples and its associated labels per characteristic. Voice samples are raw speech utterances. Every speech utterance has a duration around 3 or 4 s., but it could be different depending on the training dataset.

[0068] The labels are expressed with numerical values for characteristics, such as age, height and BMI, and categorical values for characteristics such as gender and language.

[0069] Figure 6 shows an example of the training stage for categorical values (i.e. gender) and numerical values (i.e. age). For gender (categorical value), utterances are classified in two classes: male and female. For age (numerical value), N values are used as classes. These values correspond to ages in the database, in this example ages between 15 and 90 years old are used. The same method would apply to the remaining characteristics of interest.

[0070] Throughout the training process a class-balanced database is used to properly train every class for each characteristic. The training process includes the frame selection 010, feature extraction 020, and the speaker characteristics estimation 030 modules. During the training stage, speech utterances of 3-4s are used as input to the frame selection module 010. As output, a segment of selected accumulated frames is provided. Feature extraction 020 is performed on the speech segment. The extracted feature vectors along with their labels corresponding to each speaker characteristics are used to train each DNN-based model in the speaker characteristic estimation module 030.

[0071] In one example, the DNN architecture consists of four fully-connected layers. The two hidden layers have 256 units and a hyperbolic tangent ($\tanh$) activation function each. More specifically, the $\tanh$ function is defined as follows:

$$\tanh(z) = \frac{2}{1 + e^{-2z}}$$

[0072] The output layer consists of N units and an activation function. The number of units N and the activation function in the output layer depend on each characteristic.

[0073] In case of characteristics expressed as categorical values (e.g., gender and language), the output layer consists of a layer with N units, where N is the number of classes of the characteristic of interest. For instance, for gender one can use $N = 2$ corresponding to female and male speech; for language, N is equal to the number of languages trained in the DNN. Finally, the activation function for these characteristics is a softmax function.

[0074] In case of characteristics expressed as numerical values (e.g., age, height and BMI), the output layer consists of a single unit which does the prediction. The

activation function for these characteristics is a linear function.

[0075] More specifically, the linear function is defined as follows:

$$f(z) = az + b$$

where the value of $f(z)$ increases proportionally with the value of z . The weight is represented by a , and b is the bias.

[0076] It should be noted that other DNN architectures can be used, e.g., Recurrent Neural Network (RNN) or Convolutional Neural Networks (CNNs).

[0077] Figure 7 shows an example of the test process for gender and age. Starting with selected frames in 010, a segment is progressively assembled with the accumulated frames at time t and it is processed by the live feature extraction module 020. Finally, the speaker characteristics estimation module 030 is applied.

[0078] In this example, the frames have been selected for both gender and age. Thus, speaker characteristics estimation module 030 gets a value for each characteristic for the segment of selected frames at time t . As it has been mentioned before, some characteristics are expressed as numerical values (e.g., age, height and BMI) and the other ones as categorical values (e.g., gender and language). In the example shown in the figure, the outputs are: "Female" for gender and "22.4" for age.

[0079] The speaker characteristics prediction module 030 produces estimations that are then provided to the stopping criteria module 040. The stopping criteria module 040 includes a set of methods to decide if the speaker characteristics predictions (e.g., age, height, gender, language, BMI), obtained from the speech segment of relevant accumulated frames at time t by the module 030, are reliable enough to provide an accurate estimation.

[0080] Figure 8 shows the stopping criteria 040 workflow. The stopping criteria module 040 consists of three sub-modules: minimum accumulated frames number criterion 042, weighted average estimation 043, and reliable prediction criterion 044.

[0081] Different methods could be applied for the stopping criteria, according to the type of speaker characteristics. In case of characteristics expressed as numerical values, Prediction Intervals (Pis) are used. For the categorical values, the criteria are based on the Log-likelihood ratios (LLRs). It should be noted that for both methods, Pis or LLRs are used for both the minimum accumulated frames number criterion 042 and the reliable prediction criterion 044.

[0082] For characteristics expressed as numerical values (e.g., age, height and BMI), the stopping criteria module 040 makes use of Prediction Intervals (Pis) estimated at training time. The Pis represent the variance of the characteristics prediction error; that is, the difference between the predicted and true values of the speaker characteristics. For instance, at a time t , first, the incoming

frames, that were selected as relevant for a given characteristic in 010, are evaluated by the minimum accumulated frames number criterion 042. If the criterion is satisfied, then predictions from the following selected frames are accumulated using the weighted average estimation 043. Then, the reliability of the prediction is assessed. It is important to point out that these criteria are matched with each speaker characteristic.

[0083] For the minimum accumulated frames number criterion 042, as the first speech frames may present poor speaker-characteristics information, i.e. they may not all be reliable for the estimation of speaker characteristics, one can estimate a minimum number of accumulated frames via the minimum accumulated frames criterion to avoid providing estimations which would not be accurate enough.

[0084] The minimum accumulated frames number criterion 042 is based on the prediction interval (PI) estimation. In the context of the speaker characteristics prediction task, different Pis will be obtained for each value belonging to each speaker characteristic. The prediction interval (PI) is defined as the quantification of the uncertainty on a prediction. This quantification is based on the variance of the differences between the predicted values and the true values. As a result, it will provide probabilistic upper and lower bounds that contains, with a specified degree of confidence, a future randomly selected observation from a distribution.

[0085] Before the evaluation of the minimum accumulated frames number criterion 042, the Pis must be estimated at training time. The PI estimation process consists of estimating each speaker characteristics from a training dataset, and then obtaining the Pis estimation associated to each characteristic.

[0086] Figure 9 on the top depicts the prediction interval estimation process for the minimum accumulated frames criterion for the age prediction task. The predictions, corresponding to a specific true speaker characteristic value, are used to estimate the PI for that true value. For the interval prediction (PI) estimation, one can use a training dataset S containing N speech utterances with their respective speaker characteristics information (e.g., age, height, BMI). The dataset can be defined as follows:

$$S = \{x_n, y_n^{characteristic}\}_{n=1}^N$$

where x_n denotes the speech utterances of 3-4 seconds duration and $y_n^{characteristic}$ denotes the corresponding true values or classes of the speaker characteristic (e.g., age, height, BMI, weight). The predicted characteristic values $\hat{y}_n$ obtained by the speaker characteristics estimation module 030 from the speech utterances x_n along with their respective speaker characteristics values (e.g., age, height, BMI, weight) are used to estimate the

prediction intervals (Pis). All predictions (N_c predictions) corresponding to that true value obtained by the speaker characteristics estimation module 030, are used to calculate the variance of the difference between the true value and the predicted values. The variance σ_c can be defined as follows:

$$\sigma_c = \frac{1}{N_c} \sum_{n=1}^{N_c} (y_c - \hat{y}_n)^2$$

[0087] Figure 10 shows the Standard Normal z-table for two-tailed probabilities. It represents the area or percentage under the standard normal curve corresponding to any z-score. That is, the probability of observing a z-value that is within the given $\pm$ z-score. For instance, the aforementioned probability for a $z = 1.44$ (i.e. tenths = 1.4 and hundredths = 0.04 in Figure 10), is 85.01% (i.e. $1 - 0.14987$).

[0088] Assuming that the variance of difference between the true and the predicted values σ_c shows a normal distribution, the interval of uncertainty for a specific true value can be calculated as follows:

$$interval_c = z * \frac{\sigma_c}{\sqrt{N_c}}$$

where the z value is provided by the standard normal distribution table in Figure 10. The PI around the true value y_c can be calculated as follows:

$$PI_c = y_c \pm interval_c$$

[0089] For the minimum number of accumulated frames criterion 042, the characteristic prediction from the speech segment of accumulated selected frames, obtained via the speaker characteristics estimation module 030, must be between the PI bounds at a probability determined by the z value.

[0090] Figure 9 shows the prediction interval estimation process for the age prediction task. In this example, the predicted age values had to be between the PI bounds at 85% probability, where z value is 1.44.

[0091] Once the criterion's condition is satisfied in 041 at time t, the prediction is calculated by means of a weighted average estimation 043. Thus, the following speech frames at time $t + 1$, $t + 2$, $t + 3$, ... are no longer evaluated by the minimum number of accumulated frames criterion 042 since it was already satisfied in 041.

[0092] Figure 9 on the bottom represents the evaluation of a speech segment of accumulated selected frames at a time t by the minimum accumulated frames number criterion 042 for the age prediction task. During the evaluation of the speech segment, once the predicted

age $\hat{y}_{fp}$ from the first accumulated selected speech frames at time t was obtained by means of the speaker characteristics estimation module 030, first a PI must be selected. This is performed by finding the minimum absolute difference δ between the true age y_c and the predicted age $\hat{y}_{fp}$. The difference is defined as follows:

$$\delta = |y_c - \hat{y}_{fp}|$$

Once the PI is selected, the minimum accumulated frames number criterion 042 is satisfied if the PI contains the predicted age.

[0093] For instance, for a predicted age value $\hat{y}_{fp} = 22.3$ of a speech segment, the selected PI which corresponds to a minimum δ_c is PI_{22} . Since the predicted age of the first accumulated frames $\hat{y}_{fp} = 22.3$ is between the bounds of the PI_{22} that PI is selected and the prediction is calculated by means of a weighted average estimation 043. Thus, the following speech frames at time $t + 1$, $t + 2$, $t + 3$, ... are no longer evaluated by the minimum number of accumulated frames criterion. It must be noted that z-value can be used as a parameter to constrain the prediction interval. Thus, if we desire a more conservative system, we must choose a smaller z-value.

[0094] For a predicted age value $\hat{y}_{fp} = 22.6$ of a speech segment, the selected PI which corresponds to a minimum δ_c is PI_{22} , however the predicted value is not between PI_{22} bounds. Thus, the minimum accumulated frames number criterion 042 is not satisfied, then the process continues to the frame extraction module 011.

[0095] It should be noted that there is a chance that the predicted value is contained in two PIs. For instance, a predicted age value $\hat{y}_{fp} = 21.5$ of a speech segment is contained in PI_{21} and PI_{22} . In an example embodiment, if that is the case, the final PI will correspond with the lower one, that is the PI_{21} . Other solutions could be adopted such as selecting the PI corresponding to the higher age value or the PI which corresponds to the lowest interval. Furthermore, to eliminate false alarms (meaning dealing with estimated values which has quite far from the true age value) a more conservative approach can be used by setting a smaller z-value.

[0096] In the example embodiment, a weighted average method is applied by the weighted average estimation module 043 to avoid that outlier values resulting from the estimation of the characteristics degrade the system's performance.

[0097] Figure 11 shows the system performance for the age prediction task. It is shown that the greater number of frames, the better the prediction (lower Mean Absolute Error) is and ultimately get close to the offline performance. However, there is a chance that the coming frames at $t+1$, $t+2$, $t+3$, ... could degrade the predicted value at time t due, for example, to an error in the frame selection.

[0098] To avoid this issue, one can use a weighted average making use of the frame at time t and the accu-

mulated frames at time $t-1$.

[0099] The predict value V at time t is defined as follows:

$$V_t = \beta V_{t-1} + (1 - \beta) S_t, \quad \beta \in [0,1]$$

where S_t is the predicted value at current time t and V_{t-1} is the previous weighted value at time $t-1$, and β works as a hyperparameter which weights last accumulated values and the current predicted value. Since the more frames available, the better prediction, the chosen value for β is 0.9. This way, the estimation from previous frames have a higher weight in the final decision. This approach is based on the Stochastic Gradient Descent (SDG) method with momentum used in DNNs training. Nevertheless, the β values could be tuned as necessary for each use case or application.

[0100] The reliable prediction criterion 044, is similar to the minimum number of accumulated frames criterion. It evaluates if the prediction from accumulated frames at time t , provided by the weighted average estimation 043, is reliable enough to be provided as the final estimation of the speaker characteristic with a certain confidence. If the criterion is satisfied, no more speech frames are processed and the speaker characteristic estimation is used and included in the speaker profile by means of speaker profiling module 050.

[0101] The reliable prediction criterion 044 is based on both the prediction interval (PI) estimation and the PI bounds definition at a probability determined by the z value as performed for the minimum accumulated frames number criterion 042.

[0102] Figure 12 on the top depicts the prediction interval (PI) estimation process for the reliable prediction criterion for the age prediction task. In this case, the criterion is that predicted age values must be between the PI bounds at 60% probability, where z value is 0.84.

[0103] Figure 12 on the bottom represents the evaluation of a speech segment of accumulated selected frames at a time t by the reliable prediction criterion 044 for the age prediction task. The evaluation method is similar to the one performed for the minimum accumulated frames number criterion 042.

[0104] Once the criterion's condition is satisfied, the prediction calculated by the weighted average estimation 043 is provided as a final estimation. The final result is used by the module 050 to create a speaker profile.

[0105] In case of speaker characteristics expressed as categorical values (e.g., gender, dialect, language), the stopping criteria module makes use of Log-Likelihood Ratios (LLRs), instead of PIs. LLRs are estimated at training time.

[0106] Given a segment of selected speech frames of an unknown speaker, the classification model will provide a decision whether it is likely or not that the predicted class (e.g., English or Chinese for the language speaker characteristics) can be included in the speaker profile.

For speaker characteristics expressed as categorical values, the prediction task can be restated as a basic hypothesis test between a null-hypothesis (H_0) that the predicted class does not correspond to a given speaker characteristic and an alternative hypothesis (H_1) that the predicted class corresponds to the speaker characteristic. In a mathematical form, the problem can be defined by means of a Log-Likelihood Ratio (LLR) between the likelihoods obtained for the null-hypothesis (H_0) and the alternative hypothesis (H_1), as shown in the following equation:

$$LLR_i = \text{Log}P(X|H_{i1}) - \text{Log}P(X|H_{i0})$$

$$LLR_i \begin{cases} \geq \lambda_i & \text{accept } H_{i1} \\ < \lambda_i & \text{reject } H_{i1} \end{cases}$$

where $P(X|H_{ij})$, $i = 0, 1, 2, \dots, M$ class, $j = 0, 1$, is the probability or score for the hypothesis H_{ij} evaluated for the observed X segment of selected speech frames of an unseen speaker (also referred to likelihood), and λ_i is the decision threshold for accepting or rejecting the class i hypothesis. The higher the LLR is, the greater is the confidence that it is the correct class. As shown in Figure 13. and Figure 14, for the class i and $P(X|H_{i1})$, the log-likelihood of the null-hypothesis is calculated using the equation $\text{Log}(1 - P(X|H_{i1}))$.

[0107] Figure 8 shows the stopping criteria 040 workflow, which is used for both speaker characteristics with numerical values and speaker characteristics with categorical values.

[0108] The minimum accumulated frames number criterion 042 for characteristics expressed as categorical values is based on the decision thresholds estimated from the distribution of LLRs of each class. Different decision thresholds λ_i will be obtained for each class i belonging to each speaker characteristic. These thresholds will decide if the alternative hypothesis (H_{i1}) of the class i is accepted or rejected.

[0109] Before the evaluation of the minimum accumulated frames number criterion 042, the decision thresholds must be estimated at training time. To determine the decision thresholds for each class of a specific characteristic, first calculate the LLRs from the scores obtained by the module 030 for a specific class. Then, a decision threshold for each class is determined from the LLR distributions.

[0110] Figure 13 on the top depicts the decision threshold estimation process for the minimum accumulated frames criterion for the language prediction task. Different decision thresholds are estimated for each class. From the scores obtained by the speaker characteristics estimation module 030 LLRs are calculated. From the LLR distributions corresponding to the true class and the other possible classes, the decision threshold λ is estimated. The selected threshold will determine if we want

the model to yield acceptance (score higher than the threshold) or rejection (score lower than the threshold) for every test speech segment. Two type of errors can occur: false alarm (or false acceptance), if the LLR is higher than the threshold for a non-target sample, and non-detection (or false rejection) if the LLR is lower than the threshold for a target sample. Therefore, as the LLR distributions overlap, whatever the detection threshold is, a percentage of false acceptances and false rejections will occur. The higher the threshold is, the lower the false alarms and the higher the false rejections will be. Thus, the decision threshold chosen (in most case manually) constitutes a tradeoff between the false positives and the false negative rates of the system.

[0111] For the decision threshold estimation, one can use a training dataset S containing N speech utterances with their respective speaker characteristics information, e.g., language, gender. The dataset can be defined as follows:

$$S = \{x_n, y_n^{characteristic}\}_{n=1}^N$$

where x_n denotes the speech utterances of 3-4 seconds duration and $y_n^{characteristic}$ denotes the corresponding true values of the speaker characteristic classes. The scores of predicted characteristic values $\hat{y}_n$ obtained from the module 030 and the speech utterances x_n are used to calculate LLRs. From the distributions of LLRs for each class, a decision threshold λ is set based on the tradeoff chosen between the false positive and false negative rates.

[0112] For the minimum number of accumulated frames criterion 042, one can make sure that the maximum LLR, which corresponds to LLR_i of the most likely class i , must be greater than the decision threshold λ_i , i.e. the threshold associated to the most likely class.

[0113] Figure 13 on the bottom represents the evaluation of a speech segment of accumulated selected frames at a time t by the minimum accumulated frames number criterion 042 for the language prediction task. During the evaluation of the test speech segment, LLRs are calculated from the scores obtained for each class. Then, the maximum LLR value, which corresponds to the most likely class, is selected. The maximum LLR value is compared to the decision threshold of the most likely class. The minimum accumulated frames number criterion 042 is satisfied if there is only one maximum LLR value that is greater than the decision threshold of the most likely class.

[0114] For instance, in case of the language prediction task, the predicted language class is Spanish with an LLR of 0.82. If the decision threshold for the Spanish class is 0.55, then the LLR for the Spanish class (which is the maximum LLR) is greater than the decision threshold. As the 042 condition is satisfied, the LLR is accumu-

lated by means of a weighted average estimation 043. Thus, the following speech frames at time $t + 1$, $t + 2$, $t + 3$, ... are no longer evaluated by the minimum number of accumulated frames criterion 042 since it was already satisfied in 041.

[0115] In case there is more than one maximum LLR greater than their respective threshold, the minimum number of accumulated frame criterion 042 will not be satisfied. Then the workflow will continue to the frame selection module 011 until there is only one LLR greater than its respective threshold.

[0116] The weighted average estimation 043 is similar to the one performed in the case when speaker characteristics are expressed as numerical values. There are two main differences for speaker characteristics with categorical values:

[0117] The first one is the weighted average models LLRs values instead of predicted values. The LLR value V at time t is defined as the same way as before:

$$V_t = \beta V_{t-1} + (1 - \beta) S_t, \quad \beta \in [0,1]$$

where, S_t is the LLR value for a specific class at current time t and V_{t-1} is the previous weighted average LLR value at time $t - 1$, and β works as a hyperparameter, as before. In this case, the chosen value for β is also 0.9.

[0118] The second one is that the weighted average estimation is carried out for each class of a specific characteristic. Thus, there are as many weighted average LLRs values as classes. Once the minimum accumulated frames number criterion 042 is satisfied, the weighted average LLR is estimated for each class. Then, the weighted average LLR values for all the classes are compared among each other, at time t , and the maximum value is chosen. This value is considered the maximum LLR at time t .

[0119] The reliable prediction criterion 044 for a speaker characteristic with categorical values is similar to the minimum number of accumulated frames criterion 042. It is based on both LLRs and decision thresholds as for the minimum accumulated frames number criterion 042. It evaluates if the estimated maximum LLR from accumulated frames, provided by the weighted average estimation 043 at time t , is reliable enough to provide the final class of the speaker characteristic.

[0120] The weighted average estimation 043 provides the weighted average LLR values for each class to the reliable prediction criterion 044, which is satisfied once the maximum LLR value is greater than its respective decision threshold, i.e. the decision threshold of the most likely class. The thresholds selected for the reliability criteria can be the same as the ones used for the minimum number of frames criterion or more restrictive (meaning higher thresholds could be set) as the reliability prediction criterion is the last one to be satisfied before the speaker characteristics is included in the speaker profile.

[0121] In case of more than one maximum LLR greater

than their respective threshold, the reliable prediction criterion 044 will not be satisfied. Then the workflow will continue to the frame selection module 011. If the criterion is satisfied, no more speech frames are processed and the speaker characteristic estimation is used and included in the speaker profile by means of module 050.

[0122] Figure 14 on the top depicts the decision thresholds' estimation process (which is the same process used in Figure 13 repeated here for clarity) for the reliable prediction criterion and the language prediction task. As mentioned earlier, for the reliability prediction criterion the thresholds selected can be different than the ones selected in Figure 13.

[0123] Figure 14 on the bottom represents the evaluation of a speech segment of accumulated selected frames at a time t by the reliable prediction criterion 044 for the language prediction task.

[0124] During the evaluation of the test speech segment, LLRs estimation for each class are provided by the reliable prediction module 044. Then, the maximum LLR value, which corresponds to the most likely class, is selected. The maximum LLR value is compared to the decision threshold of the most likely class. The reliable prediction criterion 044 is satisfied if the maximum LLR value is greater than the decision threshold for the most likely class. Thus, as final class the most likely class is provided.

[0125] For instance, in case of the language prediction task, the maximum LLR value from the weighted average LLR values is 1.19, corresponding to English language. Assuming the decision threshold for English is 1.00, the LLR is greater than the threshold which satisfies the reliability prediction criterion. The English language class is then used by the speaker profiling module 050 to be included in the speaker profile.

[0126] For speaker profiling 050, all the characteristics estimated from the speech signal (e.g., age, gender, height, BMI and language) are used to fill out a speaker profile in the example embodiment. Each characteristic is added to the profile progressively as soon as a reliable estimate of the characteristic has been obtained.

[0127] In summary, the novelty of the present disclosure relies on three important modules which allow the system to achieve a high accuracy with a low latency. Those modules are the DNN-based frame selection 010, the real-time speaker characteristics estimation 030, and the stopping criteria 040.

[0128] The DNN-based frame selection allows to select the relevant frames necessary to provide a reliable speaker characteristic estimation. The frame selection module 010 selects those frames that contain relevant information for estimating a given speaker characteristic. This module contributes to the accuracy and the low latency of the system.

[0129] The real-time speaker characteristics estimation allows to estimate the speaker characteristics from a speech segment of accumulated selected frames at a time t . The frame level processing contributes to the low

latency as it is not necessary to wait for the whole speech utterance to predict a speaker characteristic, in contrast with state-of-the-art systems. We estimate a speaker characteristic with only a few reliable frames.

[0130] The stopping criteria includes two different sub-modules that contribute to the accuracy and the low latency of the system: the minimum accumulated frames criterion 042 and the reliable speaker characteristic prediction 044. The first one determines if a speech segment of accumulated selected frames is accurate enough with a certain degree of reliability to be accumulated by the weighted average estimation 043. The second criterion 044 evaluates if the speaker characteristic estimation at time t is reliable enough to be provided as a final prediction that will be included in a speaker profile by means of module 050.

[0131] The techniques described herein may be implemented by one or more computer programs executed by one or more processors. The computer programs include processor-executable instructions that are stored on a non-transitory tangible computer readable medium. The computer programs may also include stored data. Non-limiting examples of the non-transitory tangible computer readable medium are nonvolatile memory, magnetic storage, and optical storage.

[0132] Some portions of the above description present the techniques described herein in terms of algorithms and symbolic representations of operations on information. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. These operations, while described functionally or logically, are understood to be implemented by computer programs. Furthermore, it has also proven convenient at times to refer to these arrangements of operations as modules or by functional names, without loss of generality.

[0133] Unless specifically stated otherwise as apparent from the above discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0134] Certain aspects of the described techniques include process steps and instructions described herein in the form of an algorithm. It should be noted that the described process steps and instructions could be embodied in software, firmware or hardware, and when embodied in software, could be downloaded to reside on and be operated from different platforms used by real time network operating systems.

[0135] The present disclosure also relates to an apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes

es, or it may comprise a computer selectively activated or reconfigured by a computer program stored on a computer readable medium that can be accessed by the computer. Such a computer program may be stored in a tangible computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, application specific integrated circuits (ASICs), or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus. Furthermore, the computers referred to in the specification may include a single processor or may be architectures employing multiple processor designs for increased computing capability.

[0136] The algorithms and operations presented herein are not inherently related to any particular computer or other apparatus. Various systems may also be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatuses to perform the required method steps. The required structure for a variety of these systems will be apparent to those of skill in the art, along with equivalent variations. In addition, the present disclosure is not described with reference to any particular programming language. It is appreciated that a variety of programming languages may be used to implement the teachings of the present disclosure as described herein.

[0137] The foregoing description of the embodiments has been provided for purposes of illustration and description. It is not intended to be exhaustive or to limit the disclosure. Individual elements or features of a particular embodiment are generally not limited to that particular embodiment, but, where applicable, are interchangeable and can be used in a selected embodiment, even if not specifically shown or described. The same may also be varied in many ways. Such variations are not to be regarded as a departure from the disclosure, and all such modifications are intended to be included within the scope of the disclosure.

Claims

1. A computer-implemented method for profiling an unknown speaker, such that a profile for a speaker includes a set of speaker characteristics, comprising:

receiving a speech signal from an unknown speaker, where the speech signal is comprised of a plurality of data frames;
for each speaker characteristic in the set of speaker characteristics, selecting a subset of data frames from a corresponding speech signal in the plurality of speech signals and thereby exclude other data frames in the corresponding speech signal from the subset of data frames,

where the data frames in the subset of data frames contain information about the speaker characteristic of interest; and
for a given speaker characteristic, generating a feature vector from the subset of selected data frames corresponding to the given speaker characteristic and predicting a value for the given speaker characteristic from the feature vector using machine learning.

2. The computer-implemented method of claim 1 wherein speaker characteristics are selected from a group consisting of age, height, body mass index, gender, accent, dialect and race.
3. The computer-implemented method of claim 1 wherein selecting a subset of data frames further comprises implementing at least one classifier for each speaker characteristic.
4. The computer-implemented method of claim 1 wherein predicting a value for the given speaker characteristic further comprises implementing a neural network for each speaker characteristic.
5. The computer-implemented method of claim 1 further comprises

assessing reliability of the predicted value for the given speaker characteristic;
adding the predicted value for the given speaker characteristic to a profile for the unknown speaker when the predicted value for the given speaker characteristic is deemed reliable.

6. The computer-implemented method of claim 5 wherein assessing reliability of the predicted value for the given speaker characteristic further comprises comparing the predicted value to a prediction interval, and deeming the predicted value reliable when the predicted value falls within the prediction interval, where the prediction interval is based on a variance of a difference between predicted values for the given speaker characteristic and true values for the given speaker characteristic.
7. The computer-implemented method of claim 5 wherein assessing reliability of the predicted value for the given speaker characteristic further comprises comparing a log-likelihood ratio to a decision threshold, and deeming the predicted value reliable when the log-likelihood ratio exceeds the decision threshold, where the log-likelihood ratio evaluates a null-hypothesis that the predicted value does not correspond to a given speaker characteristic and an alternative hypothesis that the predicted value corresponds to the given speaker characteristic.

8. A computer-implemented method for profiling an unknown speaker, such that a profile for a speaker includes a set of speaker characteristics, comprising:

a) receiving a speech signal from an unknown speaker, where the speech signal is comprised of a plurality of data frames; 5
 b) accumulating data frames from the plurality of data frames in a buffer;
 c) generating, at a given time, a feature vector from the data frames accumulated in the data buffer; 10
 d) predicting a value for a given speaker characteristic from the feature vector using machine learning, and assessing reliability of the predicted value for the given speaker characteristic; 15
 e) adding the predicted value for the given speaker characteristic to a profile for the unknown speaker when the predicted value for the given speaker characteristic is deemed reliable; 20
 and
 f) repeating steps b) - d) when the predicted value for the given speaker characteristic is deemed unreliable. 25

9. The computer-implemented method of claim 8 wherein speaker characteristics in the set of speaker characteristics are selected from a group consisting of age, height, body mass index, gender, accent, dialect and race. 30

10. The computer-implemented method of claim 8 wherein the feature vector is generated from a number of data frames which varies over time. 35

11. The computer-implemented method of claim 8 wherein assessing reliability of the predicted value for the given speaker characteristic further comprises comparing the predicted value to a prediction interval, and deeming the predicted value reliable when the predicted value falls within the prediction interval, where the prediction interval quantifies uncertainty related to the predicted value 40

12. The computer-implemented method of claim 8 wherein assessing reliability of the predicted value for the given speaker characteristic further comprises comparing the predicted value to a prediction interval, and deeming the predicted value reliable when the predicted value falls within the prediction interval, where the prediction interval is based on a variance of a difference between predicted values for the given speaker characteristic and true values for the given speaker characteristic. 45 50

13. The computer-implemented method of claim 8 wherein assessing reliability of the predicted value for the given speaker characteristic further comprises 55

es

comparing the predicted value to a prediction interval, where the prediction interval is based on a variance of a difference between predicted values for the given speaker characteristic and true values for the given speaker characteristic; computing a weighted average using the predicted value and a previous weighted average when the predicted value falls within the prediction interval, where the previous weighted average is computed using at least one predicted value for the given speaker characteristic predicted from a feature vector generated at a time earlier than the given time; comparing the weighted average to another prediction interval; and deeming the predicted value reliable when the weighted average falls within the prediction interval.

14. The computer-implemented method of claim 8 wherein assessing reliability of the predicted value for the given speaker characteristic further comprises comparing a log-likelihood ratio to a decision threshold, and deeming the predicted value reliable when the log-likelihood ratio exceeds the decision threshold, where the log-likelihood ratio evaluates a null-hypothesis that the predicted value does not correspond to a given speaker characteristic and an alternative hypothesis that the predicted value corresponds to the given speaker characteristic. 55

15. The computer-implemented method of claim 8 further comprises 35
 for each speaker characteristic in the set of speaker characteristics, predicting a value for a particular speaker characteristic from the feature vector using a corresponding neural network, where a different neural network is implemented for each speaker characteristic in the set of speaker characteristics. 40

16. The computer-implemented method of claim 8 further comprises repeating steps of accumulating data frames from the plurality of data frames in a buffer, generating a feature vector from the data frames accumulated in the data buffer, and predicting a value for each speaker characteristic in the set of speaker characteristics that is deemed unreliable, such that these steps are repeated until each speaker characteristic in the set of speaker characteristics are deemed reliable. 45 50

17. A system for profiling an unknown speaker, comprising: 55

a microphone configured to capture a speech signal from an unknown speaker;

- a frame extractor configured to receive the speech signal from the microphone and accumulate data frames from the speech signal in a buffer;
- a feature extractor configured to generate a feature vector from the data frames in the buffer;
- a speaker estimator configured to receive the feature vector from the feature extractor and, for each speaker characteristic in the set of speaker characteristics, predict a value for a speaker characteristic from the feature vector using a first set of neural networks, where the speaker estimator implements a different neural network for each speaker characteristic in the set of speaker characteristics; and
- a speaker profiler configured to receive the predicted values for each speaker characteristic from the speaker estimator, where the speaker profiler assesses the reliability of the predicted value for each speaker characteristic and adds the predicted value for a given speaker characteristic to a speaker profile when the predicted value for the given speaker characteristic is deemed reliable.
18. The system of claim 17 wherein speaker characteristics in the set of speaker characteristics are selected from a group consisting of age, height, body mass index, gender, accent, dialect and race.
19. The system of claim 17 further comprises a frame selector configured to receive the speech signal from the frame extractor, wherein, for each speaker characteristic, the frame selector selects a subset of data frames from a corresponding speech signal in the plurality of speech signals and accumulates the subset of data frames in a buffer corresponding to the speaker characteristic, thereby excluding other data frames in the corresponding speech signal from the subset of data frames, where the data frames in the subset of data frames contain information about the speaker characteristic of interest.
20. The system of claim 19 wherein the frame selector selects a subset of data frames from each of the plurality of speech signals using a second set of neural networks, where the frame selector implements a different neural network for each speaker characteristic in the set of speaker characteristics.
21. The system of claim 17 wherein the speaker profiler assesses reliability of the predicted value for a given speaker characteristic by comparing the predicted value to a prediction interval, and deeming the predicted value reliable when the predicted value falls within the prediction interval, where the prediction interval quantifies uncertainty related to the predicted value.
22. The system of claim 17 wherein the speaker profiler assesses reliability of the predicted value by comparing the predicted value to a prediction interval, where the prediction interval is based on a variance of a difference between predicted values for the given speaker characteristic and true values for the given speaker characteristic; computing a weighted average using the predicted value and a previous weighted average when the predicted value falls within the prediction interval, where the previous weighted average is computed using at least one predicted value for the given speaker characteristic predicted from a feature vector generated at a time earlier than the given time; comparing the weighted average to another prediction interval; and deeming the predicted value reliable when the weighted average falls within the prediction interval.




Figure 1




Figure 2




Figure 3




Figure 4




Figure 5




Figure 6




Figure 7




Figure 8




Figure 9

Entries in the table represent two-tailed probabilities for z statistics

	hundredths									
tenths	+0.00	+0.01	+0.02	+0.03	+0.04	+0.05	+0.06	+0.07	+0.08	+0.09
0.0	1.00000	0.99202	0.98404	0.97607	0.96809	0.96012	0.95216	0.94419	0.93624	0.92829
0.1	0.92034	0.91241	0.90448	0.89657	0.88866	0.88076	0.87288	0.86501	0.85715	0.84931
0.2	0.84148	0.83367	0.82587	0.81809	0.81033	0.80259	0.79486	0.78716	0.77948	0.77182
0.3	0.76418	0.75656	0.74897	0.74140	0.73386	0.72634	0.71885	0.71138	0.70395	0.69654
0.4	0.68916	0.68181	0.67449	0.66720	0.65994	0.65271	0.64552	0.63835	0.63123	0.62413
0.5	0.61707	0.61005	0.60306	0.59611	0.58920	0.58232	0.57548	0.56868	0.56191	0.55519
0.6	0.54851	0.54186	0.53526	0.52869	0.52217	0.51569	0.50925	0.50286	0.49650	0.49019
0.7	0.48393	0.47770	0.47153	0.46539	0.45930	0.45325	0.44725	0.44130	0.43539	0.42953
0.8	0.42371	0.41794	0.41222	0.40654	0.40091	0.39533	0.38979	0.38430	0.37886	0.37347
0.9	0.36812	0.36282	0.35757	0.35237	0.34722	0.34211	0.33706	0.33205	0.32709	0.32217
1.0	0.31731	0.31250	0.30773	0.30301	0.29834	0.29372	0.28914	0.28462	0.28014	0.27571
1.1	0.27133	0.26700	0.26271	0.25848	0.25429	0.25014	0.24605	0.24200	0.23800	0.23405
1.2	0.23014	0.22628	0.22246	0.21870	0.21498	0.21130	0.20767	0.20408	0.20054	0.19705
1.3	0.19360	0.19020	0.18683	0.18352	0.18025	0.17702	0.17383	0.17069	0.16759	0.16453
1.4	0.16151	0.15854	0.15561	0.15272	0.14987	0.14706	0.14429	0.14156	0.13887	0.13622
1.5	0.13361	0.13104	0.12851	0.12602	0.12356	0.12114	0.11876	0.11641	0.11411	0.11183
1.6	0.10960	0.10740	0.10523	0.10310	0.10101	0.09894	0.09691	0.09492	0.09296	0.09103
1.7	0.08913	0.08727	0.08543	0.08363	0.08186	0.08012	0.07841	0.07673	0.07508	0.07345
1.8	0.07186	0.07030	0.06876	0.06725	0.06577	0.06431	0.06289	0.06148	0.06011	0.05876
1.9	0.05743	0.05613	0.05486	0.05361	0.05238	0.05118	0.05000	0.04884	0.04770	0.04659
2.0	0.04550	0.04443	0.04338	0.04236	0.04135	0.04036	0.03940	0.03845	0.03753	0.03662
2.1	0.03573	0.03486	0.03401	0.03317	0.03235	0.03156	0.03077	0.03001	0.02926	0.02852
2.2	0.02781	0.02711	0.02642	0.02575	0.02509	0.02445	0.02382	0.02321	0.02261	0.02202
2.3	0.02145	0.02089	0.02034	0.01981	0.01928	0.01877	0.01827	0.01779	0.01731	0.01685
2.4	0.01640	0.01595	0.01552	0.01510	0.01469	0.01429	0.01389	0.01351	0.01314	0.01277
2.5	0.01242	0.01207	0.01174	0.01141	0.01109	0.01077	0.01047	0.01017	0.00988	0.00960
2.6	0.00932	0.00905	0.00879	0.00854	0.00829	0.00805	0.00781	0.00758	0.00736	0.00714
2.7	0.00693	0.00673	0.00653	0.00633	0.00614	0.00596	0.00578	0.00561	0.00544	0.00527
2.8	0.00511	0.00495	0.00480	0.00465	0.00451	0.00437	0.00424	0.00410	0.00398	0.00385
2.9	0.00373	0.00361	0.00350	0.00339	0.00328	0.00318	0.00308	0.00298	0.00288	0.00279
3.0	0.00270	0.00261	0.00253	0.00245	0.00237	0.00229	0.00221	0.00214	0.00207	0.00200
3.1	0.00194	0.00187	0.00181	0.00175	0.00169	0.00163	0.00158	0.00152	0.00147	0.00142
3.2	0.00137	0.00133	0.00128	0.00124	0.00120	0.00115	0.00111	0.00108	0.00104	0.00100
3.3	0.00097	0.00093	0.00090	0.00087	0.00084	0.00081	0.00078	0.00075	0.00072	0.00070
3.4	0.00067	0.00065	0.00063	0.00060	0.00058	0.00056	0.00054	0.00052	0.00050	0.00048
3.5	0.00047	0.00045	0.00043	0.00042	0.00040	0.00039	0.00037	0.00036	0.00034	0.00033
3.6	0.00032	0.00031	0.00029	0.00028	0.00027	0.00026	0.00025	0.00024	0.00023	0.00022
3.7	0.00022	0.00021	0.00020	0.00019	0.00018	0.00018	0.00017	0.00016	0.00016	0.00015
3.8	0.00014	0.00014	0.00013	0.00013	0.00012	0.00012	0.00011	0.00011	0.00010	0.00010
3.9	0.00010	0.00009	0.00009	0.00009	0.00008	0.00008	0.00008	0.00007	0.00007	0.00007
4.0	0.00006	0.00006	0.00006	0.00006	0.00005	0.00005	0.00005	0.00005	0.00005	0.00004

Figure 10




Figure 11




Figure 12




Figure 13




Figure 14

INTERNATIONAL SEARCH REPORT

International application No

PCT/ES2019/070849

A. CLASSIFICATION OF SUBJECT MATTER

INV. G10L17/26 G10L25/30
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G10L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 9 564 123 B1 (MONT-REYNAUD BERNARD [US] ET AL) 7 February 2017 (2017-02-07)	1-6, 8-12, 15-21
A	abstract; figures 1-8 columns 1-4 columns 6-8, columns 11-13	7,13,14, 22
X	W0 2016/209888 A1 (SINGH RITA [US]) 29 December 2016 (2016-12-29)	1-6, 8-12,15, 16
A	abstract; figures 1-6 paragraphs [0015] - [0060]	7,17-22

-/-



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

10 August 2020

Date of mailing of the international search report

19/08/2020

Name and mailing address of the ISA/
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Képesi, Marián

INTERNATIONAL SEARCH REPORT

International application No
PCT/ES2019/070849

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 2019/089432 A1 (UNIV NEW YORK STATE RES FOUND [US]) 9 May 2019 (2019-05-09) abstract; figures 1-5, 10-14 paragraphs [0021] - [0027] paragraphs [0080] - [0082] paragraphs [0092] - [0160] -----	1-22

1

Form PCT/ISA/210 (continuation of second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/ES2019/070849

5

10

15

20

25

30

35

40

45

50

55

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 9564123 B1	07-02-2017	US 9564123 B1	07-02-2017
		US 10311858 B1	04-06-2019
		US 2019244600 A1	08-08-2019
-----		-----	
WO 2016209888 A1	29-12-2016	US 2018190284 A1	05-07-2018
		US 2020105267 A1	02-04-2020
		WO 2016209888 A1	29-12-2016
-----		-----	
WO 2019089432 A1	09-05-2019	CA 3080399 A1	09-05-2019
		CN 111492373 A	04-08-2020
		WO 2019089432 A1	09-05-2019
-----		-----	

Form PCT/ISA/210 (patent family annex) (April 2005)